

MEMORY HUB WITH INTEGRATED NON-VOLATILE MEMORY

TECHNICAL FIELD

The present invention relates to memory systems, and more particularly, to memory modules having a memory hub and an integrated non-volatile memory for storing 5 module specific information.

BACKGROUND OF THE INVENTION

Conventional computer systems include system memory, which is typically used to store information, such as instructions of a software application to be executed by a processor, as well as data that is processed by the processor. In a typical computer 10 system, the processor communicates with the system memory through a processor bus and a memory controller. The processor issues a memory request, which includes a memory command, such as a read command, and an address designating the location from which data or instructions are to be read. The memory controller uses the command and address to generate appropriate command signals as well as row and column addresses, which are 15 applied to the system memory. In response to the commands and addresses, data are transferred between the system memory and the processor. The memory controller is often part of a system controller, which also includes bus bridge circuitry for coupling the processor bus to an expansion bus, such as a PCI bus.

Generally, the system memory of a computer system takes the form of one 20 or more memory modules that includes several integrated circuit memory devices mounted on a printed circuit board. Examples of the types of memory devices include asynchronous dynamic random access memories (“DRAMs”) and synchronous DRAMs (“SDRAMs”). Typically, the memory modules are removably plugged into connectors located on a motherboard of the computer system. The size of the computer system’s memory can be 25 increased by plugging additional memory modules into the motherboard. Memory modules are commercially available in standardized configurations, such as a single in-line memory

module (“SIMM”) and a double in-line memory module (“DIMM”), which match the connectors. The memory modules are electrically coupled to the memory controller, processor, and other devices also mounted on the mother-board using standardized memory interfaces, as well known. These standardized memory interfaces generally include a data bus, an address bus, and a control/status bus.

Often included on the printed circuit board of a memory module is a non-volatile memory in which module specific information, such as timing information, memory type, and manufacturing information, is stored. The non-volatile memory of each module can be coupled to the memory controller on the mother board through a serial bus and the connector in which the memory module is inserted. The module specific information stored in the non-volatile memory is accessed by the computer system at start-up to initialize the memory controller so that it can communicate with the memory devices of the memory module. Additionally, the basic input/output system (BIOS) or operating system of the computer system may further access the module specific information through the serial bus in performing various tasks.

A memory system that has been developed as an approach to increasing system memory bandwidth employs multiple memory devices coupled to the processor through a “memory hub.” In a memory hub architecture, or a hub-based memory subsystem, a system controller or memory controller is coupled over a high speed data link to several memory modules. Typically, the memory modules are coupled in a point-to-point or daisy chain architecture such that the memory modules are connected one to another in series. Each memory module includes a memory hub that is coupled to the corresponding high speed data links and a number of memory devices on the module, with the memory hubs efficiently routing memory requests and responses between the controller and the memory devices over the high speed data links.

A non-volatile memory is still included on the memory module for providing module specific information to the system controller of the host computer system, in the same manner as the memory module for the standard system memory

configuration previously discussed. That is, the system controller is coupled through a serial bus and module connector to the non-volatile memory in order to read the module specific information as part of initializing the computer system. With the addition of a memory hub to the memory module, a printed circuit board having more space is required.

- 5 However, in some applications, such as in hand-held computing devices or portable computers, space allocated to memory modules is at a premium, and consequently, increasing the size of the printed circuit board to accommodate the additional components is undesirable. Additionally, the time for completing initialization of the computer system upon power up will be limited by the speed at which the non-volatile memory of each of
- 10 the memory modules in a system memory can be accessed and the information transferred to the system controller over the serial bus. In applications where the demand for processing capability is immediate, minimizing the time for initializing the computer system is desirable.

SUMMARY OF THE INVENTION

15 In one aspect of the present invention, a memory hub having an integrated non-volatile memory for storing configuration information is provided. The configuration information can be copied directly from the non-volatile memory into storage registers in the memory hub. The memory hub for a hub-based memory sub-system includes a high-speed interface for receiving memory access requests, a non-volatile memory having

20 memory configuration information stored therein, and a memory controller coupled to the high-speed interface and the non-volatile memory. The memory controller includes registers into which the memory configuration information is loaded and is operable to output memory requests in response to receiving memory access requests from the high-speed interface and in accordance with the memory configuration information loaded in the

25 registers. In another aspect of the present invention, a method for initializing a memory sub-system is provided. The method includes loading configuration registers of a plurality

of memory hubs with the configuration information provided by a respective one of a plurality of embedded non-volatile memories integrated in the respective memory hub.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a partial block diagram of a computer system in which
5 embodiments of the present invention can be implemented.

Figure 2 is a partial block diagram of a alternative computer system in which
embodiments of the present invention can also be implemented.

Figure 3 is a partial block diagram of a memory module according to an
embodiment of the present invention that may be used in the computer system of Figures 1
10 or 2. .

Figure 4 is a partial block diagram of a memory hub for the memory module
of Figure 3.

DETAILED DESCRIPTION OF THE INVENTION

Embodiments of the present invention are directed to a memory hub coupled
15 to a non-volatile memory for access to information that can be copied directly from the
non-volatile memory into storage registers in the memory hub. In embodiments having the
non-volatile memory integrated with the memory hub, initialization time when powering on
a host system can be reduced since the module specific information can be copied directly
20 to configuration registers in the memory hub. Moreover, having the module specific
information copied directly to the configuration registers in the memory hub allows for a
host system to interface with the system memory without the need to accommodate any
specific characteristics of the system memory, thus, providing a more controlled
environment to which the host system may interface. Certain details are set forth below to
provide a sufficient understanding of various embodiments of the invention. However, it
25 will be clear to one skilled in the art that the invention may be practiced without these
particular details. In other instances, well-known circuits, control signals, and timing

protocols have not been shown in detail in order to avoid unnecessarily obscuring the invention.

Figure 1 illustrates a computer system 100 according to one embodiment of the present invention. The computer system 100 includes a processor 104 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 104 includes a processor bus 106 that normally includes an address bus, a control bus, and a data bus. A host bridge 110 is also coupled to the processor bus 106. The host bridge 110 is also coupled through an input/output (I/O) bus 118 to an I/O channel 120 through which one or more input and output devices can be coupled. Examples of the I/O bus 118 would be the PCI or ISA bus standards. Some common devices that would be coupled to the I/O channel 120 would be network interface cards, modems or bus adapter cards for SCSI or Fibre Channel device support. A peripheral control 124 is coupled to the I/O bus 118. Examples of peripheral control 124 devices in personal computer chipsets would be the south bridge or the I/O controller hub. The peripheral control 124 block would generally support many of the standard I/O interface functions in the system (which are not shown in the diagram) such as keyboard and mouse, which allow an operator to interface with the computer system 100. Plus common output devices such as a printer, coupled to the processor 104 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs). Peripheral control 124 would also typically be the controller or bus master for a relatively slow serial bus such as Inter-IC (I2C) or System Management Bus (SMBus) that is used by the system for housekeeping tasks such as capabilities reporting, configuration and health monitoring. The previously described components generally define a host system 101. The elements of the host system 101 are conventional, and can be implemented using designs and circuitry known by those ordinarily skilled in the art.

A system memory 132 is coupled to the host system 101, more specifically, the host bridge 110, through a high-speed bus 134. The system memory 132 is represented in Figure 1 by a memory hub based memory system that includes one or more memory modules, each of which includes a memory hub (not shown). As will be explained in more detail below, a memory hub controls access to memory devices of the memory module on which the memory hub is located. The high-speed bus 134 can be a bi-directional bus that couples together the memory hubs of the memory modules in various configurations. For example, the high-speed bus 134 can couple the memory modules together in a point-to-point configuration where information on the high-speed bus 134 must travel through the memory hubs of “upstream” memory modules to reach a “downstream” destination. It will be appreciated, however, that a high-speed link 134 providing topologies other than the point-to-point arrangement may also be used. For example, a high-speed link 134 providing a coupling arrangement in which a separate high-speed bus (not shown) is used to couple each of the memory modules of the system memory 132 to the host bridge 110 may also be used. A switching topology may also be used in which the host bridge 110 is selectively coupled to each of memory module of the system memory 132 through a switch (not shown). Other topologies that may be used will be apparent to one skilled in the art.

Additionally, the high-speed link 134 coupling the memory modules to the memory hub controller may be an electrical or optical communication path. However, other types of communications paths can be used for the high-speed link 134 as well. In the event the high-speed link 134 is implemented as an optical communication path, the optical communication path may be in the form of one or more optical fibers. In such case, the host bridge 110 and the memory modules of the system memory 132 will include an optical input/output port or separate input and output ports coupled to the optical communication path, as well known in the art.

The system memory 132 is also coupled to the peripheral control 124 through system serial busses 136. As shown in Figure 1, the system memory 132 is coupled to the peripheral control 124 through a serial bus clock line and a serial bus data

line, such as an Inter-IC (I2C) bus or a System Management Bus (SMBus), which are well known in the art. As will be explained in greater detail below, the system serial busses 136 can be used by the host system to access information from the system memory 132, such as system memory configuration information, as well known in the art.

5 Figure 2 illustrates a computer system 200 according to another embodiment of the present invention. The computer system 200 includes a host system 201 that includes the same components as the host system 101 (Figure 1). Consequently, each of the components will not be described again in detail in the interest of brevity. However, in the computer system 200, the system memory 132 is coupled to the processor 104 through 10 the high-speed bus 134. In contrast, in the computer system 100, the system memory 132 is coupled to the host bridge 110 through the high-speed bus 134. The architecture of the computer system 200 may be preferable where immediate access to the system memory 132 by the processor is desirable, such as for computer systems designed for data intensive processing applications.

15 Figure 3 shows a partial block diagram of a memory module 300 according to an embodiment of the present invention. The memory module 300 can be used in the system memory 132 (Figure 1). The memory module 300 includes a memory hub 140 coupled to several memory devices 240a-240i through a memory device bus system 150. The memory device bus system 150 normally includes a control bus, an address bus, and a 20 data bus, as known in the art. However, it will be appreciated by those ordinarily skilled in the art that other memory device bus systems, such as a bus system using a shared command/address bus, may also be used without departing from the scope of the present invention. In Figure 3, the memory devices 240a-240i are illustrated as synchronous dynamic random access memory (“SDRAM”) devices. However, memory devices other 25 than SDRAM devices may also be used. It will be further appreciated that the arrangement of the memory devices 240a-240i, and the number of memory devices can be modified without departing from the scope of the present invention.

As previously mentioned, the memory hub 140 controls access to memory devices 240a-240i of the memory module 300. Thus, memory requests and responses between the host system and the memory devices 240a-240i can be efficiently routed by the memory hub 140 over the high-speed bus 134. It will be appreciated that the system 5 memory 132 will typically include multiple memory modules, each having its own memory hub 140, which are coupled together by the high-speed bus 134. Computer systems employing this architecture can have a higher bandwidth because a host system can leverage the memory hubs 140 of the system memory 132 to access a memory device on one memory module while a memory device on another memory module is responding to a 10 prior memory access. For example, the host system can output write data to one of the memory devices in the system memory 132 while another memory device in the system memory 132 is preparing to provide read data to the processor. Moreover, this architecture also provides for easy expansion of the system memory without concern for degradation in signal quality as more memory modules are added.

15 Figure 4 illustrates a partial block diagram of a memory hub 400 according to an embodiment of the present invention. The memory hub 400 can be substituted for the memory hub 140 (Figure 3). The memory hub 400 includes a memory controller 402 coupled to a high-speed interface 404 through a memory hub bus 410. The high-speed interface 404 is coupled to the high-speed bus 134 in order for the memory controller 402 20 to communicate with the host system. The memory hub bus 410 can be implemented using conventional designs well known in the art. For example, the memory hub bus 410 can include a bus having bi-directional signal lines for receiving and transmitting signals between the memory controller 402 and the high-speed interface 404.

The high-speed interface 404 is conventional, and includes conventional 25 circuitry used for transferring data, command, and address information through the high-speed bus 134. As well known, such circuitry includes transmitter and receiver logic known in the art. It will be appreciated that those ordinarily skilled in the art have sufficient understanding to modify the high-speed interface 404 to be used with specific

types of communication paths, and that such modifications to the high-speed interface 404 can be made without departing from the scope of the present invention. For example, in the event the high-speed bus 134 to which the high-speed interface 404 is coupled is implemented using an optical communications path, the high-speed interface 404 will 5 include an optical input/output port that can convert optical signals into electrical signals for operation of the memory hub 400.

The memory controller 402 is coupled to the memory device bus 150 (Figure 3). The memory controller 402 performs the same functions as a conventional memory controller by providing control and address signals to the memory devices 240a-10 240i coupled to the memory device bus 150, and provides data signals to and receives data signals from the memory devices 240a-240i as well. However, the nature of the signals sent and received by the memory controller 402 will correspond to the nature of the signals that the memory devices 240a-240i coupled to the memory device bus 150 are adapted to send and receive. That is, the memory controller 402 is specially adapted to the memory 15 devices 240a-240i to which the memory controller 402 is coupled. More specifically, the memory controller 402 is specially adapted to provide and receive the specific signals received and generated, respectively, by the memory device 240a-240i to which it is coupled. In an alternative embodiment, the memory controller 402 is capable of operating with memory devices 240a-240i operating at different clock frequencies. As a result, the 20 memory controller 402 can isolate the processor 104 from changes that may occur at the interface between the memory hub 400 and memory devices 240a-240i coupled to the memory device bus 150, and consequently, provide a more controlled environment to which the memory devices 240a-240i may interface.

Configuration registers 403 are included in the memory controller 402. As 25 will be explained in more detail below, the configuration registers 403 are typically loaded with module specific information upon power up. The module specific information can then be used by the memory controller 402 for initialization so that it can communicate

most effectively with the memory devices of the memory module on which the memory hub 400 is located.

The memory hub 400 also includes a non-volatile memory 406 coupled to the memory controller 402 through a first configuration path 412 and further coupled to the 5 high-speed interface 404 through a second configuration path 414. As will be explained in more detail below, the non-volatile memory 406 is used to store module specific information that is used by the memory controller 402 during initialization. The non-volatile memory 406 can be implemented using conventional non-volatile memory, such as FLASH memory or other types of electrically erasable programmable read-only memory 10 (EEPROM). The non-volatile memory 406 is preferably embedded memory formed as part of the memory hub 400, and can be of a relatively small capacity, such as 256 Kbits or 512 Kbits. However, other types of non-volatile memory, and different capacities can be used as well without departing from the scope of the present invention.

The memory controller 402, high-speed interface 404, and non-volatile 15 memory 406 are also coupled to a local system serial bus 420. The local system serial bus 420 can be coupled to a host system through a system serial bus, such as the system serial bus 136 shown in Figures 1 and 2. The non-volatile memory 406 is used to store information specific to the memory module on which the memory hub 400 is located. Examples of the module specific information includes timing information for the memory 20 devices of the memory module, memory module configuration data, memory device type, manufacturer data, and the like.

As previously mentioned, in conventional memory modules, the module specific information is typically accessed by a host system upon start-up to properly initialize the host memory controller so that it can communicate most effectively with the 25 memory devices of the memory module. In contrast, however, the non-volatile memory 406 is integrated with the memory hub 400 so that the module specific information can be accessed and copied directly from the non-volatile memory 406 to appropriate configuration registers 403 in the memory controller 402 when the host system is powered

on. In embodiments having the non-volatile memory 406 integrated with the memory hub 400, and having the module specific information copied directly to the configuration registers 403, initialization time when powering on a host system can be reduced. Moreover, having the module specific information copied directly to the configuration registers 403 of the memory hub 400 allows for a host system to interface with the memory module without the need for the host system to accommodate any specific characteristics of the memory module or the memory devices on the memory module. Thus, the host system can interact with a system memory generically, relying on the memory hub 400 to manage the specifics of the memory module.

The non-volatile memory 406 can also store module specific information used by a host system as well, such as memory module capacity, memory module clock speed, and the like. Such information is often used by the basic input/output system (BIOS), the operating system, or application software in performing various tasks. For module specific information that should be provided to the host system in which the memory module is located, the information can be provided through the high-speed interface 404 to the host system via the configuration path 414 and the high-speed bus 134. Alternatively, in embodiments of the present invention having the local system serial bus 420, the module specific information can be provided though a system serial bus that is coupled to the local system serial bus 420.

From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.